

NETWORK MANAGEMENT SYSTEM

Inventors

Elango Gannesan

Taqi Hasan

Allen B. Rochkind

Sagar Golla

Express Mail No. EL 901895764 US

Prepared By

VIERRA MAGEN MARCUS HARMON & DENIRO LLP

0900482 070601
T09070 28400660

NETWORK MANAGEMENT SYSTEM

Inventors

Elango Gannesan
Taqi Hasan
Allen B. Rochkind
Sagar Golla

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention is directed to management of network services in a data center, and in particular to management of services, subscribers, devices, log servers, and facilities using a common, virtualized management system.

Description of the Related Art

Public wide area networks such as the Internet have expanded the types of services used and demanded by enterprises of their network infrastructure. As the number, complexity and interaction of the services has risen, the associated costs of both the infrastructure itself and maintaining the infrastructure have risen as well. Many enterprises have turned to outsourced vendors, sometimes called a managed service provider or a data center, to provide these services in lieu of building and maintaining the infrastructure themselves. Customers of such managed service providers are called subscribers.

The managed service provider can operate in many different ways. Typically it can provide secure facilities where the infrastructure service

equipment is located, and manage equipment for the subscriber. The scope of management and services is defined by an agreement with the customer calling for the managed service provider to solely or jointly manage the equipment with the subscriber. In other cases, the managed service provider can lease the physical space from another provider (called a hosting provider) and provide just the management of the infrastructure equipment on behalf of its subscribers.

A data center is a specialized facility that houses Web sites and provides data serving and other services for subscribers. The data center may contain a network operations center (NOC), which is a restricted access area containing automated systems that constantly monitor server activity, Web traffic, and network performance and report even very slight irregularities to administrators so that they can spot potential problems before they happen. A data center in its most simple form may consist of a single facility that hosts all of the infrastructure equipment. However, a more sophisticated data center is normally an organization spread throughout the world with subscriber support equipment located in various physical hosting facilities.

Data centers allow enterprises to provide a number of different types of services, including e-commerce services to customers; extranets and secure VPNs to employees and customers; firewall protection and Network Address Translation (NAT) services, web caching and load balancing services, as well as many others. These services can all be provided at an off-site facility in the data center without requiring the enterprise to maintain the facility itself.

The equipment that provides the infrastructure services for a set of subscribers can take several forms, depending on the implementation. Depending on the complexity and variety of services required, the equipment generally includes one or more single function devices dedicated to the subscriber. Generally, because the devices are designed with the co-

location model in mind, service devices generally include the ability to provide only one or a small number of services via the device. Generally, typical multi-function devices combine services that are closely related, such as NAT and firewall services. A data center facility generally has a number of devices to manage, and in many cases the devices multiply as redundant devices may be used for fail over security to provide fault-tolerance or for load balancing.

Figure 1 shows a typical single facility data center 20 and exemplary network architecture within the data center facility 20. It should be recognized that Figure 1 is oversimplified for the purpose of showing the configuration of how such a data center facility is typically managed, and there are numerous additional components and devices in a data center facility not shown in Figure 1. As shown therein, in one configuration, each subscriber has a leased physical cage – a cabinet of hardware which may include service provision devices and the subscriber's application servers as well as other specialized equipment for implementing the subscriber's service structure.

As shown therein, the data center facility 20 is coupled to a Wide Area Network (WAN) 50 via a high-speed interface device, such as an ATM switch 55. It will be recognized that the particular type of physical network to which the data center is coupled is merely illustrative and not germane to the presentation of the invention.

As illustrated in Figure 1, the WAN may be a worldwide system of computer networks such as the Internet. Each ATM switch may be coupled to one or more level 2 and level 3 OCI layer switches 60, which direct traffic to any number of subscriber cages 22. Each subscriber cage may be leased by a particular subscriber of the data center, and may include equipment dedicated to servicing a particular subscriber. The subscriber may be, for example, an Internet business or company which seeks to offload its network operations to the data center. As shown in Figure 1, each subscriber cage

includes equipment from a different subscriber – “E-Client1.com”, “E-Client2.com”, etc. It should be understood that various forms of service models between the subscriber and the data center have been developed, and the particular form of business arrangement of a leased cage – sometimes referred to as “co-located” servers is illustrative only. Each cage is a physical rack of appliances dedicated to the particular subscriber and may be coupled to a network back plane and the application servers, which are maintained and configured by the subscriber in conjunction with the network operations center.

In general, in the typical data center case, the subscribers service equipment is designed with one subscriber in mind and hence, a data center providing outsourced management services to many subscribers must provide a separate set of infrastructure equipment for each subscriber. This equipment can come from many different vendors. The cages may include network appliances dedicated to one or more of the following tasks: routing, firewall, network address translation (NAT), SSL acceleration, virtual private networking, public key infrastructure (PKI), load balancing, Web caching, or the like. As a result, the management of all subscribers within the data center becomes very complex and expensive with many different management interfaces for all of the subscribers and subscriber devices. Administering the equipment in each cage is generally accomplished via an administrative access interface coupled to each single function device.

The difficulty in administering a data center as shown in Figure 1 is that changes to each subscriber’s individual configuration must be made at the cage, or at minimum, via appliance specific configuration mechanisms. Normally, subscribers themselves have no control over the service appliances and the data center administrators manage these appliances. As will be readily apparent, the more subscribers one has in the data center, the more resources must be committed to administration. This can become exceedingly difficult where changes must physically occur at the cage or via

individual devices, especially in multi-facility data centers. In general, the data center will include one or more network operations centers, and one or more facilities operations centers. The network operations center generally refers to the facility which manages all physical facilities and the data center.

5 The data center may have one or more physical facilities, each of which has its own facilities administrator who can have physical access to each of the cages. Hence, there can be at least two types of data center administrators depending on the organization of the data center. If the data center has all of its infrastructure equipment within a single facility, then there are only the
10 personnel of that single facility acting as data center administrators. In larger data centers spread throughout the world, there is normally a central headquarters of the data center provider along with numerous separate data center facilities throughout the world. At the headquarters the data center may provide the capability to manage or monitor any device throughout the
15 world that the data center is using to provide to its subscribers infrastructure services.

In order to alleviate some of these management problems, some virtual solutions have been developed, allowing administrators to configure devices via network interfaces. However, such devices do not administer
20 services on an object level. That is, they do not allow the administrator access to administrative functions on a service level, but rather allow administrators to administer one or more particular types of devices via the device interface, by providing a common connection point for a number of devices in the data center. One example of such a solution is provided by
25 Arula Systems corp. These devices allow an administrator to connect to a service device and through the administration device, manage a multitude of service devices.

Hence, such solutions are of limited scalability and scope.

SUMMARY OF THE INVENTION

The invention, roughly described, is a virtual management system for a network facility, such as a data center, or any facility having a plurality of components which can be organized as objects for presentation in a virtualized environment. Once organized into the virtualized management system of the present invention, management of the data center is streamlined and less user intensive.

Various aspects of the invention are described herein. In one aspect, the invention comprises a virtual management system for a data center, and includes a management topology presenting devices, facilities, subscribers, log servers, and services as objects to an administrative interface; and a configuration manager implementing changes to objects in the topology responsive to configuration input from an administrator via the administrative interface.

In an exemplary embodiment, the user interface is a graphical user interface designed to work in a platform independent environment.

In a further aspect, the invention comprises a management interface for a network. The management interface includes a graphical user interface presenting a plurality of network items as objects within the interface; service applications coupled to the graphical user interface objects, the applications controlling configuration of network objects responsive to the user interface; and a network manager interacting with devices on the network to implement changes provided by the service applications.

In yet another aspect, the invention is a graphical network interface for a data center. The graphical network interface includes a plurality of object views, including a facility object view, a subscriber object view, a device object view, a log server object view, and a services object view. Each said view includes a set of objects organized by a hierarchy relative to another of said views. The graphical network interface also includes at least one link to

an object in said set of objects, allowing modification of configuration data for the object via the view.

In a further embodiment, the invention comprises a system for providing centralized management to a data center having at least one facility and a plurality of devices configurable to provide network services to subscribers. In this embodiment, the invention includes a management server coupled to the plurality of devices; and an interface to the management server including a configuration interface allowing a user to configure each of the network services provided by the devices in the data center. In one aspect, the management server communicates with the devices, downloading configuration data to and uploading configuration data from, the devices. The management server and the interface may communicate via a LAN, WAN or the Internet.

In a still further embodiment, the invention comprises a multi-facility management system wherein a management server is provided in a first facility, a set of devices to be managed is provided in a second facility, and the management server communicates with the devices via a WAN or the Internet between the facilities via a secure protocol.

In another aspect, a management hierarchy regulating access to objects in a virtual management system is provided. The hierarchy includes a data center administrator having at least access to configure services for all subscribers, facilities and devices in the data center; a facilities administrator having at least access to configure services for subscribers and devices at a particular geographic or virtual facility; and a subscriber administrator having at least access to configure all devices assigned to perform tasks for the subscriber.

The present invention can be accomplished using hardware, software, or a combination of both hardware and software. The software used for the present invention is stored on one or more processor readable storage media including hard disk drives, CD-ROMs, DVDs, optical disks, floppy disks, tape

drives, RAM, ROM or other suitable storage devices. In alternative embodiments, some or all of the software can be replaced by dedicated hardware including custom integrated circuits, gate arrays, FPGAs, PLDs, and special purpose computers.

5 The advantages of the present invention will appear more clearly from the following description in which the preferred embodiment of the invention has been set forth in conjunction with the drawings.

10 In the following detailed description, the present invention is described by using block diagrams to describe either the structure or the processing that implements the method of the present invention. Using this manner to present the present invention should not be construed as limiting of its scope. The present invention contemplates both methods and systems for implementing a network management system. In one embodiment, the system and method of the invention can be implemented on general-purpose
15 computers. The currently disclosed system architecture may also be implemented with a number of special purpose systems.

20 Embodiments within the scope of the present invention also include articles of manufacture comprising program storage apparatus and having encoded therein program code. Such program storage apparatus can be any available media which can be accessed by a general purpose or special purpose computer. By way of example, and not limitation, such program storage apparatus can comprise RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired
25 program code and which can be accessed by a general purpose or special purpose computer. Combinations of any of the above are also included within the scope of such program storage apparatus.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention will be described with respect to the particular embodiments thereof. Other objects, features, and advantages of the invention will become apparent with reference to the specification and drawings in which:

Figure 1 depicts typical network architecture within a data center facility, including a subscriber cage having single function network appliances.

Figure 2 depicts management of a data center facility using the virtual management system of the present invention.

Figure 3 depicts a management topology in accordance with the present invention.

Figure 4 depicts a system architecture implemented in accordance with one embodiment of the present invention.

Figure 5 is a depiction of system architecture for a multi-facility environment.

Figures 6A- 6E are depictions of graphical user interface views utilized in the system of the present invention.

Figure 7 is a depiction of the modules running on the network management server, and the management server agent for interacting with service appliance devices in one embodiment of the system of the present invention.

Figure 8 is a depiction of the administrative management hierarchy utilized in accordance with one embodiment of the system of the invention

DETAILED DESCRIPTION

A virtualized data center management solution is disclosed herein. The invention is presented in multiple aspects and embodiments. In one

aspect, the invention disclosed is a centralized management administration system with multiple subscribers. In this embodiment, each subscriber is managed as if the infrastructure equipment (each set of devices) were dedicated to an individual subscriber. In a further embodiment, all subscribers share common infrastructure equipment, and the management system segregates the management information to ensure security. The virtualized network management system allows multiple subscribers of a data center to be managed within a single integrated system.

In general, the management system of the present invention may be understood as an organization of data objects referred to as the virtualization topology, shown in Figure 2. The virtualization is normally made visible to the management administrator via some management application, and the topology structure is supported by an appropriately configured database. Within the virtualization topology, objects are comprised of devices, subscribers, facilities, log servers and data centers.

Figure 3 shows a network management architecture utilized with a virtual management system in accordance with the present invention. As shown therein, a single administrative access interface 75, such as a stand-alone administrative application running in a platform independent process, may be used to administer a multitude of services provided to subscribers. In the example shown in Figure 3, multi-function service devices 100 capable of providing a multitude of services to the subscribers are managed by administrative access interface 75. As shown in Figure 3, administrative access interface 75a provides access to the devices 100 via the Internet, or to the devices directly within the data center either within the physical facility via the secure network within the data center facility itself (75b). The multifunction devices 100 are coupled to the application servers and the network within each subscriber cage. The administrative access point may comprise a graphical user interface (GUI) or a command line interface (CLI), both of which are accessible through a multitude of applications.

It should be recognized that the virtual management system of the present invention need not be used with multifunction service devices 100, but could alternatively be used with devices providing a single or small number of functions. As will be described with respect to the system architecture set forth below in Figure 4, when utilized with such multi-service appliances, the virtual management system of the present invention provides particular advantages in performance and administration. A multifunction appliance suitable for use with the virtual management system of the present invention is disclosed in co-pending United States patent application serial number _____ entitled CONTENT SERVICE AGGREGATION SYSTEM, filed July 6, 2001, [Attorney Docket No. NEXSI-01025).

In the virtualized management system of the invention, a single system administrator can configure individual devices, log servers, or services for a subscriber, the facility, or for the entire data center, all via the interface 75. In the multifunction devices described with respect to co-pending application serial number [NEXSI-01025US0], the internal management approach is to segregate individual subscribers within the multifunction device to prohibit crossover services and avoid security problems within the device itself.

An exemplary system architecture for implementing the virtual management system of the present invention is shown in Figure 4. In this architecture, the multifunction appliance 100 is shown as being connected via a network to a network management server (NMS) 200. A secure network protocol HTTPS's may be used as the means of communication between the NMS and the multifunction appliances. The NMS can provide a graphical user interface to the administrative interface device 75 to allow a system administrator to configure appliances. A network graphical user interface 210 may comprise an application developed in a platform independent code to run in a browser application, such as Netscape 6.0 or Internet Explorer 4. or greater, 5.0 or greater, or a platform specific administration application

running on a host device. It will be understood that the GUI may be delivered to the administrative interface device by any number of methods, including servlets, a network transmitted self-installation package, or provided on a machine-readable medium. All such embodiments of the administrative interface communicate with the NMS to provide changes to devices, services, administrative policies or the like, as described below.

In the example shown in Figure 4, the GUI is a Java based graphical user interface that is provided by the management server via the network to a web browser using java servlets. The NMS GUI includes management applications such as a policy editor, a status editor, and a log/alert viewer.

Logging and monitoring servers 230 are provided in the data center to record events and traps from the multifunction appliance and provide standard format output reports to the administrator.

In a further aspect of the present invention, also described in further detail below, a subscriber may be allowed to administer all the services that are accorded to that particular subscriber by the data center. This further reduces the burden on the network operations center of configuring particular services for the subscriber.

In general, the NMS server 200 communicates with other components of the system via some secure protocol (e.g. HTTPs). (It should be recognized that within the data center, a secure network may be configured between the various components coupled to the NMS, so that HTTPs need not be required.) Network management server 200 receives and stores policies and configuration settings which are input using the user interface. The NMS 200 further includes policy databases and regulates access controls, network object definitions and security policies, as well as log server settings and reports. Using the GUI, an administrator can configure device and service configuration data and policy information. This information is stored in the network management server 200 and downloaded securely to each of the multifunction devices 100. In addition, the management server

may include a network management back-plane application (or device agent) to allow for management applications to couple to the multi-function appliances 100. The NMS server can also perform certificate management, version control, and enforce access control to the devices.

5 In the foregoing description, the invention is described with respect to an embodiment wherein certain objects of the management system are described in detail as implemented using particular protocols or integrated into web-browser applications. It should be understood, however, that in the present invention, each component of the system - server, client, data
10 center, facility, etc. represents an object and the invention is implemented to support the Common Object Request Broker Architecture (CORBA).

Administrative Interface

The GUI acts as a configuration input mechanism for the virtual
15 management system of the present invention. Although the system will be described with respect to a particular implementation using a GUI, it will be understood that the functions performed by the GUI may be implemented using a command line interface (CLI) as well. In addition, it will be recognized that the NMS and the GUI may communicate via any direct, LAN, or WAN
20 connection, or communicate via a collection of networks, such as the Internet.

Any of the topology objects, data center, subscribers, facilities, devices, or log servers can have new instances created via the GUI. Thus, for example, a new subscriber topology object is created when a new
25 subscriber is added to the data center. Any number of copies of the GUI may be in operation at the subscriber company's administrative center, to manage that subscriber's portion of the data configuration policy.

Figure 5 shows a multi-facility, multi-subscriber embodiment of the system architecture of the virtual management system of the present
30 invention. As shown therein, a series of subscribers, 300, 310, each of which

has its own graphical user interface, couples via a WAN 60 to a network operations center 350. Network operations center 350 includes its own version of the graphical user interface 320 for access by the NOC or facilities administrator who may be physically present in the NOC. All three interfaces
5 are coupled to the network operations center's network management server 200a. The data center network operations center facility 350 may include its own subscriber equipment, but in the example shown in Figure 5, separate data center facilities 360 and 370 each include a plurality of subscriber cages 368,369, 378,379, and multifunction devices 362, 364, 372, 374. In data
10 facility 360, two multifunction devices 362, 364, coupled in a redundant fashion allow for failsafe rollover between the devices in the event there is a hardware failure in one of the devices. Each of the devices 362, 364 is coupled to one or more subscriber cages 368, 369, as well as a log server 366. A similar configuration exists in data center facility 370. Facilities 360
15 and 370 are coupled via a WAN 70 to the network operations center 350. In this embodiment, a single network management server 200a is utilized. It should be recognized that the network management server 200 may actually comprise one or more physical servers configured as a single virtual management server. It should be further understood that the WAN's 60, 70
20 may comprise public WANs, secure networks, or a series of interconnected networks, such as the Internet.

Using the GUI, the management administrator perceives the virtualization topology through a set of views. As used herein, a "view" is a hierarchical layout of the data center, its facilities, subscriber, devices and/or
25 log servers. A view is normally shown as the familiar tree structure, with the root being the data center, and the branches being the facilities, subscribers, devices and log servers. It will be recognized that other GUI representations other than the tree structure may be incorporated into the system of the present invention without departing from the scope and
30 character of the invention.

Each view can have varying depths showing levels of granularity in the data center configuration. When an object is selected, various management capabilities can be performed on the object. The properties of the object can be viewed or edited. Each object like data center, facility, subscriber, device, or log server has a set of properties that can be managed via the interface. One example of an object property is the name property. Other properties are dependent on the type of management object. For example, a subscriber object may have a list of infrastructure services that the subscriber has contracted with the data center. As discussed in further detail below, the system incorporates a privilege-based model of administrator access, with each level of privilege providing a more fine grain access of what services can be managed. As such, views and levels of views may be prohibited to certain levels of administrators based on that administrative user's policy definition.

Exemplary GUI views are shown in Figures 6a – 6e.

Figures 6a – 6d are views available to global or NOC administrator. Figure 6a shows a global data center view sorted by facilities/devices and subscribers. At this level, an administrator can view all facilities, the devices shown in those facilities (as shown in Figure 6a), and the subscribers served in those facilities. In addition to viewing devices and services in the facility, this view allows definition of application access control privileges, the definition to the application of the data center NOC and facilities, specification of what multi-function devices reside in what data center facility, and policy database attributes. This level allows the administrator to define such things as descriptive and naming information for NOC and the facilities. Other device specifications can include mappings of device IP addresses to devices in a specific facility, or mappings of IP addresses to the various subscribers and subscriber service bundles. This view may also indicate how many devices reside in each data center facility, and how these devices will be used for fail over or load balancing.

5

10

15

20

25

30

service bundle associated with it. The policies and configurations for a specific service bundle configured for the subscriber are specified on a service-by-service basis (e.g. routing, firewall, NAT, VPN, PKI). If a public key certificate represents a subscriber's service bundle, then PKI for that service bundle is configured here.

5

Numerous other types of views are possible. For example the following two level views are possible:

ROOT	BRANCH	BRANCH
Data Center	Subscribers	Facilities (Fig. 6c), or devices
Data center	Facilities	Subscribers (Fig. 6b), devices (Fig. 6a), or log servers (Fig. 6d)
Data Center	Log Servers	Subscribers or Devices
Data center	Devices	Subscribers or log servers

10 The following three level views are possible:

ROOT	BRANCH	BRANCH	BRANCH
Data Center	Subscribers	Facilities	Devices
Data Center	Facilities	Subscribers	Devices
Data center	Facilities	Devices	Subscribers
Data center	Facilities	Devices	Log Servers
Data center	Facilities	Log Servers	Devices
Data center	Log Servers	Devices	Subscribers

As will be readily apparent, numerous levels and varieties of object views are possible. Views with a depth of four levels are possible. For example, under the data center we list the facilities. For each facility, we list the log servers. For each log server, we list the devices that can log to the log server. At the fourth level for each device we list those subscribers on the device that can log to the specified log server.

It should be further recognized that in the aforementioned examples, the root level may comprise the facility (to a facility administrator), or a subscriber (to a subscriber administrator), rather than the data center view (shown to the NOC administrator). The root of the tree view can depend on the level of access granted to the administrator utilizing the application.

Some examples of how the user interface may be used are as follows. When the application is started for the first time, configuration data needs to be entered by the NOC administrator. A selection of a particular item in the tree view – for example by double clicking on a subscriber, can bring up the Read/Modify dialog for the subscriber. Similar operations are used for facilities, devices and services. This allows the selected object's definition information to be viewed in more detail and if access privileges permit, that information can be changed.

If one of the applications in the "Applications" menu is clicked, then that application is brought up for the particular object. E.g. if the "Services" application is clicked in the menu, then a subscriber's service configuration/policy editor is brought up. Views may be changed using the view menu command structure

In some cases, it may not make sense to show views and submenus. If a facility is selected, the "Services" application menu option is not accessible (being subscriber oriented only). However, the "Logging" and "Status" applications are accessible, giving logging, alerting, and statistics information for the facility.

The applications menu contains the non-administrative applications: "Services", "Logging", "Status". The "Services" item is active only if a subscriber or device is the current selected object. In that case, clicking on "Services" brings up the service policy/configuration application. "Logging" brings up the log event application, which reports log events and alerts within the scope of the object selected in the component tree pane. E.g. if a subscriber is selected in the component tree pane, then log events/alerts for only that subscriber are displayed by the application. If a data center facility is selected, then the logging application presents all events/alerts coming only from that data center facility. The "Status" application displays statistics tables (and possibly dynamic graphs of selected statistics parameters) within the scope of the object selected in the component tree pane.

For example, the subscriber administrator may only have the facilities by subscriber view except only the subscriber's own hosting data center facilities are shown. The "Subscribers By Facility" and "Devices By Facility" views are not shown.

25

Each network management server provides a number of service modules which may include, for example, a client connection manager module 250, a policy manager module 252, an authentication module 254, a topology subsystem 256, a certificate manager 258, a Servlet engine 261, an HTTP client 262, a web server 264, a request handler 266, a performance monitor 268 and an alert/trap monitor 270. Each of these service modules communicates with a network management back plane application 260 in the device to configure device functions and report on device operations. It should be recognized that while the service modules are illustrated in the context of being a part of the NMS server, all or a portion of the modules may be incorporated into the user command interface, such as the GUI, which is uploaded to the administrative interface client. This allows for the distribution of the processing load of the virtual management system. In addition, other service modules may be incorporated into the management server, as needed allowing for maximum flexibility in scaling the virtual management system as new services, devices or functions are required.

The network management server 200 may also serve as a content services application server, hosting a plurality of content services applications. As shown in Figure 7, these content services applications can include a OSPF/RIP router 285, Network address Translation (NAT) 282, Firewall 284, VPN/IPSEC with IKE and PKI 286, and Bandwidth Management/QOS 288. It will be recognized that other service applications may be provided as necessary given the types of services required by the subscriber and provided by the service devices 100 in the data center.

The system management of infrastructure services (like firewall, NAT, VPN, QOS, web caching, web load balancing, SSL acceleration, etc.) for multiple subscribers is normally performed via an infrastructure service applications or Content Services Applications. Because all of these infrastructure services for all of the subscribers are running on a single virtualized management system, each content services application can be

centrally launched for each selected subscriber. When an application is launched for that subscriber, the application runs only in that subscriber context. That is, any application such as, for example, the "services application, is created only for that subscriber and all configuration data pertaining to that subscriber is available in only that subscriber's context. The management information (which could be configuration or policy information) generated by this application applies only to that subscriber.

The virtual management system incorporates a content management server (in this case the Network management server) which hosts the management services for subscribers supported by a data center. These services are accessible to subscribers through a single external URL. The NMS can transparently switch to service that was specific to a subscriber. There could be several log servers that are used by a subscriber, at a given time, in a given deployment configuration. Based on the context, the NMS directs the incoming connection to the appropriate log/web server securely. These server process resources are transparent to the end-user. This process is called Intelligent Switching and results in ease of use, higher performance and better scalability, allowing the data center to scale with subscribers, as the demand for hosting services grows.

Other service applications can be run for each subscriber. For example, a virtualized management system may supply a log view reporting application for each subscriber. The various logged events from the log servers are reported to this application, but only for the selected subscriber. Another service application could be a status reporting application. Alerts generated on behalf of the subscriber can be displayed in this status service application. Similarly network or security management statistics can be displayed for that subscriber by the status viewing service application. If an infrastructure service like VPN or SSL acceleration requires digital certificates, then the certificate management application (described below)

can be launched for each subscriber to provide the certificates required by the infrastructure service.

In addition to launching service applications in the subscriber context, other service applications may be launched in a device context. For example the physical configuration of the device (e.g. setting port addresses) can be performed for each device in the data center. This device configuration application is centrally launched for the selected device virtualization topology object. The log view and status view service applications might also be launched in a device context. They would display logs and status information that emanated only from the selected device.

Similarly, service applications may be launched in a log server context. For example, the log viewer application launched in this context would show log events coming only from the selected log server.

Still further, service applications can be launched in the global data center context. For example certificate management in the data center context could be used to generate certificates for all the components of the network management system to provide secure network communication for network management transactions. The log view and status view service applications can be launched in the data center context and log events and status information from every device in the data center could be displayed.

A service application can be launched in any number of topology object contexts. Nevertheless, it would not make sense to launch certain types of applications in certain contexts (such as, for example, launching a certificate management application in a facility context). The service applications that can be launched may not necessarily be those mentioned explicitly above; it can be any application that makes sense in the specified topology object context

The management service modules interact with the content services applications to provide the aforementioned service management. The Client Connection Manager 250 regulates the number of connections between

management client and network management server 200. It handles such actions as: sign on requests; requests for subscriber policies; requests for policy installations; requests to create an administrator; requests for device information; requests to change device information in the network management server database; requests for VPN policy information; requests to create new subscribers; requests to change subscriber information and requests to add new device information to the network management server database. A request/response mechanism is used via the embedded Web Server and use a secure protocol (e.g. HTTP) to establish a session to initiate different requests.

The policy management application 252 enforces subscriber level access control policies to maintain the data integrity and validates policies installed. Before the rules are installed, they are checked for integrity and for any possible security holes in a rule base, as well as to determine whether there is any conflict in any of the other service application rules. The policy management application can retrieve rules on a subscriber basis. This application also handles definition of VPN policies and command line interface definitions of router configurations.

The authentication manager 254 provides access control and regulated administrative privileges. Access to the module may be password encrypted and implements the security scheme set forth below.

The topology subsystem 256 maintains the relationship between objects in the virtual management system. Topology information is stored as a table in topology database file.

The certificate manager 258 is a centralized public key infrastructure (PKI) manager for each subscriber. It interacts with multiple certificate authorities and their databases for enrollment and certificate download/renewal and key management.

The servlet engine 261 provides mapping of administrator comments into servlets and interfaces to access http structures to which configuration results transmitted to the devices 100 can be written.

5 The client 262 is utilized with the back plane application 260 to communicate with the devices and supports GET and POST http methods required for the back plane, as set forth below. It should be recognized that the invention described herein could be utilized with non-HTTP based protocol clients, and non-web-browser enabled clients. ****

10 The embedded web server 264 is utilized to respond to HTTP requests from browsers to establish communication sessions with administrative interfaces 75 which may be, for example, utilizing a web browser to retrieve the GUI to administer the data center.

15 The request handler 266 is implemented as servlets which store data to respond to requests from the application modules. For example the policy handler would serve a policy request based on the access control context established during the logon. Policy information is stored in XML format and the request handler enforces concurrency control and uses the servlet engine to store data.

20 The performance monitor 268 is supported through the query interface on predefined attributes. One may also set up a polling interval to query attribute values through the GUI.

25 The alert/traps module 270 stores system critical information received by the NMS server via a secure protocol from the devices 100. The module then updates an alert viewer in the user interface in real time. The alert viewer can also be used to fetch this information from the NMS using HTTP requests.

30 Each service provided by the multifunction service devices include manageable elements represented by data structures tailored to the functional and performance requirements of that service. The management back plane 260 maps service specific data structures and the transfer

syntaxes supported for network management interactions between the device and the management server. In one embodiment, a "management back plane" tool such as Rapid Logic's RapidControl™ Backplane available from Wind River Systems, Inc. is used to perform those mappings.

5 Alternatively, all manageable elements may be mapped to a SNMP management information base.

The network management back plane application 260 may include facilities for communicating with the network components of the invention including an SNMP agent 262, a CLI parser/engine 264, a Telnet engine 268,

10 a network management engine 272, a policy configuration engine 276, and an embedded Web server 278. Each of the back plane components converts communications from the service applications of the network management server 200 to the communication required for the multifunction service appliances 100, and interprets responses from the multifunction appliances

15 100 and logging servers 230. The management back plane further allows for third party management and reporting applications to interact with the network management server to present information on configuration of the system as well as logging ports via a third party standard formats.

The SNMP agent 242 is primarily used for monitoring the overall

20 health and basic functionality of the multifunction appliance through a third party network management application. (One example of such an application is HP Openview available from the Hewlett-Packard Company). Various counters and statistics are supported for each service enabled in addition to relevant management information base data. SNMP Traps will also be

25 supported for the Network Management application.

The CLI/Parser Engine 244 is called when input to the device is received through a telnet session, or by a serial input through an out-of-band port on the multifunction device 100. Once a request comes in, a lookup is performed on the request and matched to a database linking the requests to

30 an operations code in one of the manager applications. When the CLI Engine

receives the data, it packages it back as a response to the appropriate application session.

5 The telnet server 246 allows command line interface interaction with the management server and mimics the same functionality as a serial connection to the multifunction device, set forth above.

The log server engine establishes a TCP/IP session with the log server to continuously send logging events to the log server through a secure channel. Alerts are also sent to the logging server through this channel.

10 Data aggregation engine 274 provides for data aggregation through an IPC mechanism in the multifunction device and is responsible for sending requests to various services in the device for statistics and general data collected at each service. Essentially the IPC does a lookup and dispatches requests based on the results of this lookup.

15 The policy configuration engine 276 uses a policy string or set of strings that mimics well-known or commercially utilized policies for services such as firewalls which are sent to the multifunction device for VPN and Firewall policies. The engine takes incoming configuration, parses it and hands it off to the IPC mechanism. The configuration request and parameters are then sent to the appropriate service (which will configure its individual parameters).

20 The embedded web server is used in communicating to the Network Management Server. The device receives configuration information through HTTP over SSL in XML format. Again, these are parsed by the Policy Configuration Engine and dispatched across the multifunction device.

25 The network management back plane is a component of the back plane application used for managing the device 100 and server 200 interaction mappings. In the embodiment utilizing the aforementioned commercial back plane management application, the system includes markers and pointers to allow the application to efficiently control system interactions. It should be understood that some form of internal application

30

09900482 09900482

management structure may be utilized in accordance with the present invention, and the particular type of management structure is not relevant to the scope and content of the present invention.

To increase the usability of the network management system a set of reporting facilities is provided. One type of reporting is to make available logs for various events detected by the infrastructure service equipment. Such are, for example, the detection of "anomalous" conditions where the types of anomalous conditions are configured), detection of security threats, congestion notification, alerts, etc. These logged events are stored on the log servers 230. The log servers accumulate the raw event information reported by the infrastructure devices, and reporting applications have access to the log servers so that the raw information can be converted into useful reports. Normally because the logging function is traffic intensive, a data center may support multiple log servers to distribute the logging load. Many different logging distribution schemes are possible. A typical log server distribution scheme is to provide one or more log servers within a single data center facility and have the devices only within that facility log to those log servers. It should be recognized that a plurality of different types of log server configurations within the data center are possible while remaining within the scope and content of the present invention.

Administrative Management Hierarchy

In a further unique aspect of the invention, the virtual management system includes an administrative hierarchy allowing different levels of system administrator access to varying levels of configuration and inspection.

Figure 8 shows an overview of the administrative management model utilized in the context of the virtual management system of the present invention. The users of the virtualized network management system are called administrators. The administrators that manage these services can be either personnel of the data center or personnel of the subscriber company

which contracts with the data center or both. The role that the data center or subscriber administrator plays in managing these services is determined by contractual agreement between subscriber and data center. Thus the virtualized network management system supports both data center and subscriber administrators.

As noted above, there are numerous types of data center administrators, including facility administrators and NOC administrators.

Facility administrators are located at individual facilities of the data center administrative throughout the world. Facility administrators generally have the closest contact with the subscribers that have infrastructure services managed by that facility. Subscriber administrators are personnel that work for the subscriber customer of the data center that have contracted with the data center to manage and/or monitor the management operations pertaining only to that subscriber.

Each type of administrator is given a set of management capabilities. These capabilities are described in terms of management scope and access controls.

The types of objects and the properties of those objects that can be managed by a type of administrator define the management scope of the administrator.

Subscriber administrators are limited to monitoring or modifying management information that pertains only to that subscriber. Although the virtualized management system contains management data for all subscribers, each subscriber administrator manages only his/her part of the management database. The subscriber administrator has no access (neither read or write) to any other subscriber's part of the database.

Since all of the infrastructure equipment, the devices and log servers belong to the data center; the subscriber administrator is prohibited from managing any of the corresponding device and log server topology objects. Also any attributes that pertain to the data center as a whole or to any data

center facility are off-limits to the subscriber administrator. The subscriber administrator has potential capability to manage or monitor only the infrastructure services that the subscriber had contracted for with the data center or to receive logging or status reports generated by the infrastructure service equipment.

The facility administrator is limited to the management of only those objects that pertain to his/her data center facility. This means that the facility administrator can manage only the devices and log servers of their data center facility. The facility administrator can manage the infrastructure services of only those subscribers that are hosted on the devices within his/her data center facility. The facility administrator can manage attributes that pertain only to their own facility. For example, the facility administrator can add or delete devices and log servers from the facility. The facility administrator can add or remove subscribers that have services to be hosted only at that facility.

As used in the present model, the NOC administrator has no scope limitations. The NOC administrator can potentially manage or monitor any device, log server, or subscriber service in the entire data center. This capability is potential because it is limited only by the access controls specified for this type of administrator as discussed below.

The administration model allows multiple administrators of any type to concurrently manage the system with each type of administrator limited to their scope. The access may be further limited by any concurrency controls that may be in effect in order to prevent conflicts that corrupt the management database.

Access controls may be assigned to particular administrators. There are many access control schemes however there are some basic concepts that apply to establishing any access control scheme for administrators in the virtualized network management system.

The access control scheme will specify whether the administrator has no access, read access only, or read and write access to any specific part of the management database within the scope of a particular administrator. Write access means that the administrator can make changes to the specified part of the database.

If a specific part of the management database is available in scope to a data center or subscriber administrator, the access controls granted to either the data center or subscriber administrator should be determined by agreement made between subscriber and data center when the subscriber had contracted with the data center. For example if the subscriber had granted the data center exclusive right to configure or modify the infrastructure services contracted for, then only the data center administrators will have write access to these services.

Administrative access controls determine which administrators can create new administrator accounts or can modify or delete existing administrator accounts. Subscriber administrators granted these special privileges can only create, modify, or remove subscriber administrator accounts of administrative personnel working only for that subscriber. Facility administrators having these privileges can manage accounts for administrators only of that facility or of subscribers hosted by that facility. NOC administrators having these privileges can manage administrators from any facility or any subscriber company.

Administrative access controls may determine which data center administrators can modify the virtualization topology within the scope of that administrator. Such privileges give such administrators the right to add or remove devices or log servers or change their properties.

A large number of access control scheme variations are possible. For example write access can be granted to anything within the scope of the administrator. Or write access may be granted only for modifying services for a set of subscribers. Or write access may be granted only for certain

services, but not to other services. All such schemes may be incorporated into the virtual management system of the present invention.

Integrated Services

5 The management system of the present invention allows for an selective services management. That is, all services subscribed to by the user may come under the control of the virtual management system, and subscribers need only contract for those services that are required. For example, subscribers that require only firewall services, but not VPN services
10 need not contract for such services and can save the associated costs of implementing the VPN service.

 In addition the centralized nature of the services allows integration of services. Integrated services management allows changes made to one service having an effect on other services provided to the subscriber to
15 immediately propagate to the other services of the subscriber. For example, if a user requires a VPN and that user has also contracted for firewall and NAT services, a change to add the VPN via the virtual management system is integrated with changes to the firewall and NAT servers though the changes made to the VPN application.

20 More specifically, in implementing a VPN, access though a packet filtering firewall is required. In addition, static mappings in a NAT service allowing users to point a VPN client at a given address for access to the servers of the subscriber are required. In this example, a subscriber level administrator or higher can configure the parameters of the VPN via the VPN
25 service application, by setting for example, the type of authentication used and the IP address of the VPN server, and the VPN service application will communicate with the routing, NAT and firewall applications to map the static IP, allow access to a certain port, such as port 25, for IPSec VPN traffic validation, and thereby enable the VPN for the subscriber while minimizing

In one embodiment, this is implemented using each content services application. The content services application for the VPN, for example, will search through rules to allow appropriate settings for protocols, such as IPSec and IKE, to pass through a firewall implementation, and change appropriate DNS settings, both of which are implemented by other systems to implement a VPN. Similarly, the VPN application may make changes to NAT settings implemented by the NAT service application

Yet other variations include the presentation of the management interface to system administrators. Command line interfaces are supported, including interfaces which mimic commands from popular third party configurations of companies who make single function service provision products. In addition, the GUI itself may have a different configuration than the Tree-style configuration set forth herein, and the manner in which the GUI is supported on the management interface – be it a non-specific port of platform independent code, a platform specific embodiment, or an application designed to run in a host browser – are all within the scope and context of the present invention.

Attorney Docket No.: NEXSI-01025US0
Z:\nexsi\1025\1025.app.doc

to limit the invention to the precise form disclosed. Many modifications and variations are possible in light of the above teaching. The described embodiments were chosen in order to best explain the principles of the invention and its practical application to thereby enable others skilled in the art to best utilize the invention in various embodiments and with various modifications as are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the claims appended hereto.

5

09900482 070604